Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Intel® XScale™ Microarchitecture Serves Up Breakthrough I/O

Lance Packer
Senior Product Marketing Engineer
I/O Products Division
Intel Corporation

Overview

I/O bandwidth is becoming a critical issue in the design of Fibre Channel-based Storage Area Network (SAN) equipment, Network Attached Storage (NAS) disk arrays, high-performance RAID controllers, and RAID-on-motherboard (ROMB) solutions in servers.

Breakthrough I/O performance for these and other intelligent networking applications becomes a reality with the announcement of the Intel® 80310 I/O processor chipset with Intel® XScale™ microarchitecture. The Intel 80310 processor chipset includes two chips—the new high-performance Intel® 80200 processor based on Intel XScale microarchitecture, and the Intel® 80312 I/O companion chip, which provides system level support functions.

The convergence of voice, data, and video and the adoption of hardware-based RAID to protect critical data are creating a voracious appetite for increased I/O speed and bandwidth. Intel is meeting the demand with the Intel® Internet Exchange™ Architecture, a flexible framework for adding intelligence to network infrastructure and communications appliances.

Intel's new I/O processor plays an important role within the Intel Internet Exchange Architecture. With a clock speed of 733 MHz at less than 1.3W and combined with highly integrated peripherals, the 80310 I/O processor chipset can deliver 300 percent more I/O operations/second in RAID applications, this compared with Intel's previous I/O processors. See Figure 1.

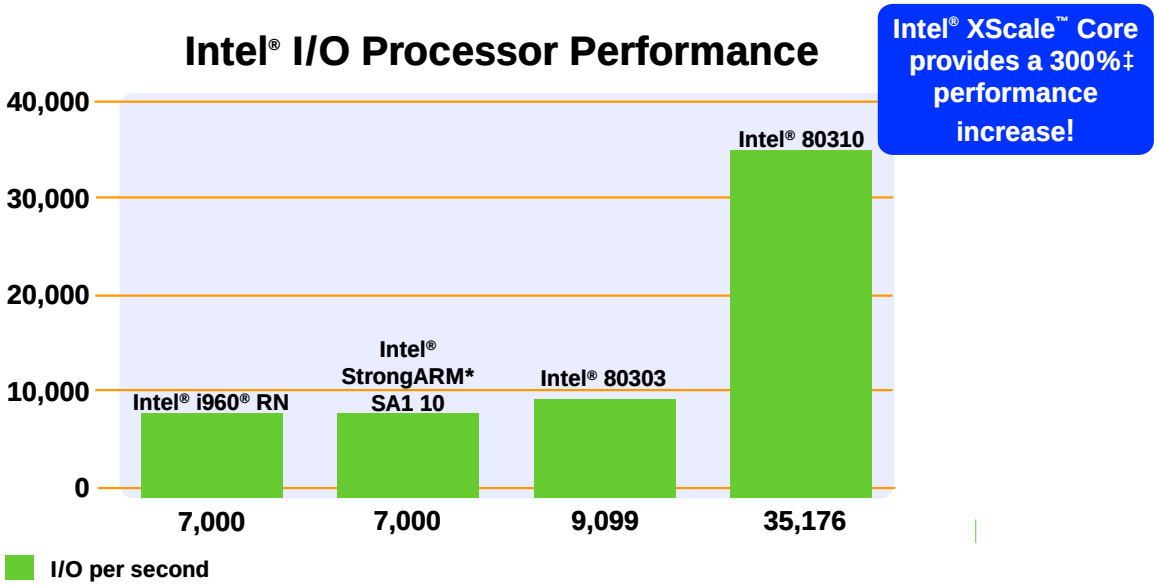

Notes: Intel performance testing. Performance will vary depending on many design factors. 2-KB block size, random I/O, database, Web server, or file server workloads.

‡Estimated performance. Performance may vary depending upon a variety of factors.

Intel® XScale™ Microarchitecture

The Intel XScale microarchitecture offers low-power features ranging from one ten-thousandth of a watt (in standby mode) to 1.3 watts, and performance capabilities allowing operation at clock speeds from zero (in standby mode) up to 733 MHz. The Intel 80310 I/O processor chipset will be released with multiple clock speed bins up to 733 MHz. The Intel XScale microarchitecture core is manufactured on Intel's advanced 0.18-micron process technology. See Figure 2.

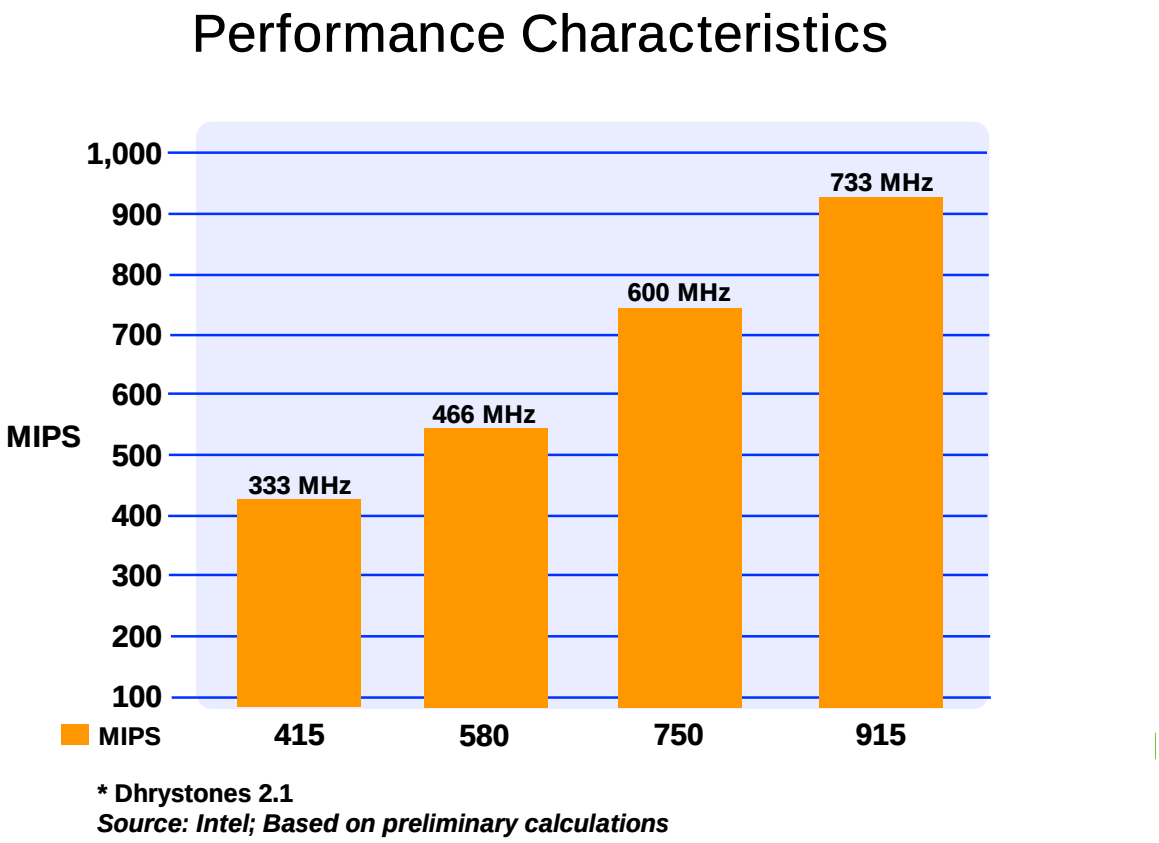

*** Dhrystones 2.1**
Source: Intel; Based on preliminary calculations

In addition to a seven-fold increase in core speed, the 80310 chipset features Intel® Superpipelined RISC technology, combining high speed with extremely low power requirements.

The new microarchitecture features a number of throughput enhancements, including:

- An expanded 7-stage integer pipeline and 8-stage memory pipeline—implemented for faster performance
- Dynamic branch prediction and extensive data bypassing
- Preservation of the ARM* v.5TE instruction set compatibility for increased code density
- Data and instruction cache size increases to 32 KB
- A new 2-Kbyte mini data cache added to accelerate video and audio streaming.

Easing the Transition

Millions of Intel® i960® I/O processors have been shipped to-date, and the Intel 80310 I/O chipset provides a solid transition path for the many developers who are already familiar with Intel i960 processor technology. The 80312 I/O companion chip, one of the two chips of the fully validated Intel 80310 I/O processor chipset, supplies the I/O peripheral logics; its implementation can speed time-to-market, because it shares identical I/O mapping and peripheral compatibility with the i960 I/O processor series. The Intel 80310 chipset is peripheral set compatible with the recently announced Intel® 80303 I/O processor.

Intel's previous I/O processors are known for tight integration of I/O peripherals, and the Intel 80310 I/O chipset is no exception. The Intel 80312 I/O companion chip features an integrated 66-MHz, 64-bit PCI-PCI bridge, a 100-MHz, 64-bit SDRAM memory controller supporting up to 512 Mbytes of ECC memory, an enhanced 1-Kbyte Application Accelerator Unit to support hardware XOR for RAID data parity calculations, six secondary PCI clocks, universal (3.3V and 5.0V) PCI support, eight general-purpose I/O interfaces, and flash memory interface. See Figure 3.

Intel® 80310 Chipset Block Diagram

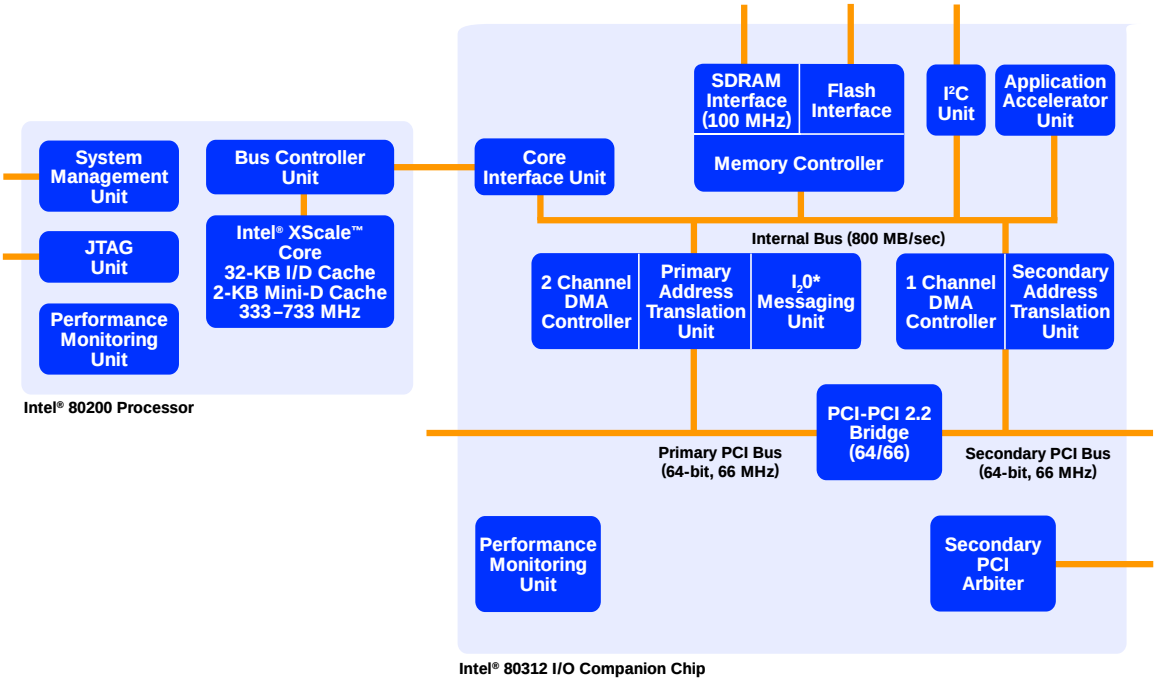

The new chipset also provides advantages for developers who are already conversant with Intel's StrongARM* architecture. The 80310 chipset is code-compatible with the Intel® StrongARM SA-110 chipset and is compliant with ARM architecture v.5TE.

To speed product development, the 80310 I/O chipset is supported by a comprehensive development environment including the GNU* and ARM tool chains and an extensive list of third-party vendors.

Processor-intensive Applications

Now undergoing evaluation by more than 100 Intel customers, the first product implementations of the 80310 I/O chipset will likely be high-performance SCSI and Fibre Channel RAID controllers.

Another example, SAN, relieves network bottlenecks by connecting multiple hosts to centralized storage arrays over dedicated channels based on Fibre Channel, SCSI, or other high-bandwidth I/O technologies. Because it centralizes storage, SAN places a premium on I/O throughput.

Other applications include RAID implementations in servers, in addition to Internet infrastructure devices built with Intel Internet Exchange architecture to support processor-intensive data encryption, security algorithms, and video and audio data streams. The 80310 chipset has the capability to handle extremely complex next-generation applications in many of these categories, while providing the first step on a roadmap to higher performance I/O processors.

Summary

With the convergence of voice, data, and video data streams supported by Intel Internet Exchange architecture, I/O bandwidth has never been more significant in the design of network infrastructure devices and communications appliances. The proliferation of high-performance data storage solutions such as Storage Area Networks, Network Attached Storage, and Intel® Integrated RAID controllers, makes I/O throughput a critical determinant of network performance.

Intel's I/O Products Division is meeting the challenge with the breakthrough performance of the new 80310 I/O processor chipset, the first product to be introduced with Intel XScale microarchitecture. It's a solution that combines 733-MHz clock speed with tight integration of I/O peripherals, compatibility with prior-generation I/O processors, and a large and growing choice of development tools. It all adds up to breakthrough performance for developers of the new connected world.

More Info

For more information on the Intel 80310 I/O processor chipset, visit the I/O Building Blocks page on the Intel Developer Web site.

For more information on Intel® processors with Intel XScale microarchitecture visit the Intel XScale microarchitecture area of the Intel Developer Web site.

Author Bio

Lance Packer is a senior product marketing engineer in Intel's I/O Products Division, where he is responsible for development, launch, and support of Intel's family of I/O processors. Lance joined Intel in 1998, bringing over 12 years of industry experience with him. He holds a B.S.E.E. from Utah State University.

Columns

From the Editor
Donna Loveland
Managing Editor
Intel Developer Update Magazine
Intel Corporation

Column

News from last month's Intel® Developer Forum Conference continues to emerge, with Intel's new Xscale™ microarchitecture, free remote development services for IA-64 applications, and a new wireless client architecture topping the list. This issue gives you a closer look at an innovative approach to using S3 sleep state technology showcased in San Jose, plus updates on advances in PXE remote booting and 1394b implementation.

Intel® XScale™ Microarchitecture Serves Up Breakthrough I/O—cover story—Breakthrough I/O performance for networking applications becomes a reality with the announcement of the Intel® 80310 I/O processor chipset, the first product with Intel® XScale™ microarchitecture.

Remote Access to Pre-Release IA-64 Systems—Intel® Early Access Services offers software developers cost-free, maintenance-free, high-speed Internet access to fully functional, secure IA-64 systems.

Accelerating Internet Expansion to Wireless Clients—The Intel® Personal Internet Client Architecture separates hardware and software design environments so applications can be developed independently, eliminating serial development and speeding time-to-market.

Deskpro IAPC Design Example—Compaq's Deskpro EN series demonstrates a variety of ways OEMs can use S3 technology to differentiate their products from other platforms.

Real-Time 1394b Data Transfer for Consumer Electronics—By choosing the right implementation of IEEE 1394b, manufacturers can achieve maximum throughput for the lowest cost and still provide significant product differentiation.

PXE Manageability Technology for EFI—PXE (Preboot Execution Environment) provides standardized remote installation and manageability for enterprise network clients, and now it's available for EFI (Extended Firmware Interface).

For details about ongoing developments, stay tuned to Intel Developer Update—your conference between the conferences.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

Departments

Applied Computing

Accelerating Internet Expansion to Wireless

Randall Smith
Director of Group Marketing
Intel Wireless Communications and Computing Group
Intel Corporation

Overview

Advances in communication technology, as well as the pervasiveness of the Internet are fueling the development of a new generation of mobile, Internet-enabled, wireless clients, and applications. In the Internet-everywhere world of today, mobile users are demanding that their wireless clients have more functionality and performance to handle everything from enhanced computing, to communications, to entertainment applications. Today's consumers are demanding a broad range of wireless clients, including personal digital assistants (PDAs), smart phones, Web tablets, e-books, and other mobile clients.

A number of technological innovations are driving the evolution of the next generation of wireless clients to user expectations:

- Third-generation cellular standards, which provide more bandwidth for new data-intensive software applications that integrate Internet and multimedia capabilities.
- Emergence of low-power, high-performance microprocessors, dense memory, and efficient baseband logic.
- Low-cost, high-performance servers that address the interface between information sources and wireless clients, and which make end-to-end capabilities a reality.
- Distributed communications technologies enabled by service-discovery software middleware, such as JINI, CORBA*, COM/DCOM, JAIN, and others.

These emerging technologies are putting intense demands on the data processing and data management capabilities of wireless clients.

Redesign of Mobile and Handheld Platforms

The coming wave of next-generation wireless standards and customer requirements is already creating a need to redesign today's mobile and handheld computing platforms. It is also creating a need to change the current serial development process. Because of this push for change, many of the design standards that worked in the past for cellular phone and handheld clients now also require redesign.

The standards must be enhanced or changed to deliver optimized support for new, sophisticated services enabled by new, third-generation wireless bandwidth. Also, as the industry rushes to develop next-generation wireless clients and services, a new model for parallel hardware and application development must emerge. At its foundation, the parallel development process requires an open, flexible, scalable, and extensible base architecture.

Parallel HW and SW Development

New products and services must be developed at a rapid rate to keep pace with the emerging market. This represents a challenge for the wireless industry because of limitations in today's serial product-development practices, applications development processes, and network operator system-validation practices.

Today's application-development and product-introduction process for wireless devices is not only serial, it is unwieldy and slow. Hardware (the silicon and device) must be developed first. Applications are then written for the target hardware and air interface, and this means that developers must be very involved in the communications stack software. Finally, both hardware and applications must be tested and receive Type Approval from each country in which the devices will be used. Currently, every application must be designed and tested specifically for each air interface and network.

This kind of development puts a huge burden on developers to go through the approval process for each product. And, with technology advancing so rapidly, applications cannot keep up with the current serial development and delivery process. One way to speed up the development process is to decouple hardware and software development. When the two design environments are separated, hardware, computing applications and communications applications can be developed independently, thus saving valuable development time.

The industry requires a new client and application development process if it is to supply the broad range on wireless Internet-enabled products and services that mobile consumers will demand. Intel believes that to keep pace with the Internet, development of hardware and applications must be decoupled and executed in parallel, with the applications written to a general-purpose processor.

In parallel development environments, applications are developed independent of the communications stack. This minimizes the need for network operator testing and validation. A parallel development process will allow hardware and applications development to evolve at their own speeds. The Intel® Personal Internet Client Architecture (Intel® PCA) provides the new parallel development environment. This environment will allow a new community of hardware and application developers to emerge, resulting in a broad range of new Internet-enabled hardware and software products.

Intel's Wireless Solution

Today's rapidly advancing technologies require not only parallel hardware and applications development, but also a wireless architecture that deals with both current and future capabilities. These capabilities include enhanced multimedia functions, color animation, M-commerce, connectivity, communications, battery life, location-based applications, and proximity-use. While many companies provide individual components for specific wireless products, Intel PCA provides the underlying structure and platform for the wireless Internet client. Intel's expertise in PC (personal computer) architecture, microprocessors, and networking and Internet infrastructure give the company a unique and effective perspective on Internet architecture in wireless space.

Intel PCA meets this requirement by providing an optimum framework for the development of next-generation wireless Internet clients and applications. The new architecture will redefine the performance and functionality expectations of both hardware and application developers.

The architecture has many exciting attributes including:

- Open software interfaces between major subsystems, for easy integration and expansion by third parties
- Scalability through the architecture's major building blocks. By adding or interchanging the building blocks, developers can create a broad range of possible configurations
- High level of modularity, allowing each module to be independently tested and reused across many different systems
- Extensibility for future uses not originally defined, made possible by a high degree of abstraction in the original design
- A modular design philosophy to simplify Type Approval testing when new applications are added to the platform
- Operation-system (OS) savvy for supporting numerous real-time and application OSes
- Support for all major cellular air interface standards via pluggable protocol logic and software components
- Ability to incorporate security mechanisms into e-Commerce applications

The architecture will provide greater freedom for both hardware developers and applications developers to implement their designs, as well as a broad range of solutions. The client software and hardware can be implemented with the wireless modem or as separate applications. The Intel PCA interfaces allow modem functions to be changed and/or upgraded for the air interface without requiring a rewrite of the application or the OS environment. In summary, the Intel Personal Internet Client Architecture facilitates a high degree of reuse and promotes a modular design.

General Benefits

The Intel initiative separates the hardware and software design environments so that they can be developed independently. Basically, Intel is providing a common architecture that can be used across many platforms and in many configurations, from low-power to high-end functionality.

A common architecture has another major benefit: there is already a large installed base of potential customers. Last year, there were 444 million cellular subscribers, of which 5 million had Internet access. By 2004, the estimated number of cellular subscribers is expected to be 1.4 billion, while the number with Internet access will be over 560 million. This is considerable financial incentive to use a common architecture on which to base increasingly complex applications.

Software Developer Benefits

Software developers will find that a common architecture provides these additional benefits:

- Eliminates the need to wait for hardware approval before being able to develop software.
- Rapidly accelerates time-to-market for products.
- Lets manufacturers put their software applications on a broader range of platforms, from PDAs to smart phones to automotive clients.
- Reaches a global customer base with more software more quickly.

Hardware Developer Benefits

Hardware developers will find that a common architecture has these major advantages:

- The availability of broad-range applications allows developers to more easily differentiate products.
- Since hardware vendors try to address different hardware segments, the common architecture and general-purpose processors ensure that products can be compatible with many types of wireless Internet applications.
- The Intel PCA architecture is designed to address almost all major interfaces so that developers have a global (multiple-market), product offering.
- The architecture is scalable, supporting a wide range of processing power and power usage. This means that the processor can handle very low power platforms as well as high-end, high-drain performance.

With the Intel PCA wireless architecture, hardware developers can design a single platform that can be used over a broad range of products (cell phones, smart phones, PDAs, automotive clients). They can improve their time-to-market and, at the same time, become more efficient in their use of engineering resources.

Summary

Intel is focused on accelerating Internet expansion to wireless clients. This is a clear statement of Intel's commitment to bringing the hardware and software elements of the Internet together through a common architecture for wireless clients. Specifically, the Intel Personal Internet Client Architecture will assist application developers by providing universal hardware and software interfaces for development of wireless handheld devices.

Key members of the wireless community are already working to take advantage of this architecture. Developers who want to accelerate their wireless client product development process should consider basing their next wireless client application on the Intel Personal Internet Client Architecture.

More Info

For more information about the Intel PCA initiative, refer to Intel's white paper, located on the Intel Developer's Web site. Developers can also find additional information about wireless technology at that site.

Author Bio

Randall Smith recently joined Intel as director of marketing for Intel Wireless Communications and Computing Group, where he will be focusing on marketing Intel's vision, technology, and product capabilities to the wireless Internet industry. In the past, Randall has worked for Samsung Electronics as director of marketing for their U.S. cellular products, and at Motorola in the cellular marketing group. Randall has more than 20 years experience in the wireless and communications markets.

Desktop

Deskpro IAPC Design Example

Louis Hobson
Senior Member, Technical Staff
Commercial Desktops Division
Compaq Computer Corporation

Overview

The energy-efficient, S3 Suspend-to-RAM sleep-state technology is relatively new. As with other new technologies, a careful implementation of S3 allows original equipment manufacturers (OEMs) to maximize the benefits of the technology in desktop PC platforms.

Compaq's Deskpro EN series is a clear example of how OEMs can use S3 technology to differentiate their products from other platforms. With an eye toward real-world compatibility issues, Compaq has engineered special features into their S3-compliant Deskpro EN series of personal computers (PC). This allows Compaq to offer a more robust system implementing S3 technology.

What Is S3?

S3 is a sleep state for PCs in which power is supplied only to essential components, such as system memory and wake-capable devices. In older technology (the S1 sleep state), devices such as the graphics card and various PCI devices do not lose power. However, in the S3 sleep state, to radically reduce power consumption in modern PCs, peripheral and other devices do lose power. When a PC is placed in the S3 sleep state, volatile system context is saved to memory, and system power consumption drops (ideally) to 5 watts or less. Trickle power is then maintained in communication subsystems and peripherals, such as modems, and in wake-capable buses (such as PCI and USB). When the PC is awakened, the previous state is restored from memory as full power resumes.

PCs in the S3 sleep state must be able to quickly return to a fully functional state when triggered by a wake event, such as keystroke, mouse movement, telephone ring (wake on ring, or WOR), LAN (wake on LAN, or WOL), and ping. This allows the PC to return to functionality in time to answer the phone, receive a FAX, and send or receive e-mail.

The Compaq Approach

While much of the hardware and BIOS design of a Suspend to RAM (STR) or S3 platform is fairly static (stable or fixed), there are still many opportunities to differentiate one PC from another. The features Compaq chose for S3 compliance not only make their Deskpro EN systems more robust for this new technology, but add many distinct features to this Instantly Available PC reference platform.

For example, Compaq identified several design issues in S3 implementations. To deal with these issues, Compaq engineered extra features into their BIOS, such as hard disk reset during POST, video re-POST enable, and special support for ATAPI devices. These features gives the Deskpro EN a solid S3 implementation, with the added benefit that the system can also handle most instances of various OSes, peripherals, and add-in cards.

Along with other features, the Deskpro EN series includes these important differentiators:

- Custom power supply
- Video re-POST BIOS setup options
- Int13 hard-disk reset
- USB device power, which supports wake from USB devices including keyboards and mice
- PS/2 mouse and keyboard wake up
- Ability to enable/disable wake from mouse as a user-selectable BIOS option

The Compaq message to other developers is that, despite being a new technology, S3 provides significant benefits to the user. There are many features you can add to your systems that will make S1-to-S3 design transition easier and still provide important product differentiation.

Custom Power Supply

For some time, platform OEMs have been asking OS manufacturers for power budgeting features for wake-capable devices. With the advent of S3 standards and technology, power budgeting features are even more important. However, OEMs want to provide power information in some sort of table that indicates the kind of power the power supply can provide in S3. In contrast, OS manufacturers want the OS to read this information directly from the power supply, an activity that requires a more expensive implementation of the power supply.

Without the correct electrical information, the PC may allow too many wake devices to be enabled, a situation in which the power required in S3 exceeds the standby power of the power supply. Because in today's systems, the power information cannot be read directly from the power supply, OS manufacturers have continued to omit power budgeting features from their OSes. This makes power budgeting features an after-market application or the responsibility of the platform OEM. (Intel offers power budgeting software as an after-market application)

There are several solutions to working around the current lack of power budgeting features. One solution is to increase the capabilities of the power supply. Another solution is to restrict the number of PCI wake-capable devices allowed in the system. The disadvantage to the first approach is cost: power supplies that are more robust are also more expensive. In the second approach, restricting the number of wake-capable devices can be a nuisance to customers who want to add more devices than the system can support.

On the Deskpro EN, Compaq wanted to make sure that users could install as many wake-capable devices as there were plugs, so it designed a custom power supply with increased auxiliary power. Specifically, the Compaq power supply provides 2 amps at 3.3 volts and 1.7 amps at 5 volts. That's enough power for a wake device in every one of the Deskpro EN's four PCI slots, plus plenty of extra power for all USB wake devices. And the Deskpro EN still draws far less than the current Environmental Protection Agency ENERGY STAR requirements.

With Compaq's custom-designed power supply, users don't have to worry about whether there is enough power in the sleep state to support all the devices that need to receive that power. Are there too many USB devices on the system? Will the system lose both main and aux power? These are no longer issues with Compaq's robust, custom-designed power supply.

Re-POST Video Option ROM

The Deskpro EN provides several BIOS setup options that help make the S3 experience more enjoyable. These options help resolve issues raised by the variety of video cards available for today's systems, and with the customer penchant for installing older, non-S3-compliant cards in S3-compatible systems.

Enabling the re-POST video option ROM feature tells the BIOS to re-execute the graphics option ROM upon S3 resumption. If the STR-capable video card is well behaved, users should never have to enable this feature.

Like other product differentiators, this feature can be implemented in a variety of ways. However, in general, a flag is set that tells BIOS to re-POST video and re-execute the option ROM. This puts the noncompliant video card back into a state that the video card driver can handle. Context is saved and restored by the BIOS, since the noncompliant driver (or video card) cannot do that itself.

When a video card is not compatible with S3, customers can use the video re-POST feature until a more compliant driver or video card is installed. Almost all noncompliant video-card drivers accept the re-POST procedure because POST doesn't affect the default actions of the drivers; the drivers were not going to restore context anyway.

Int13 Hard Disk Reset

Compaq also provides the Int13 hard-disk reset as a BIOS setup option. Compaq provides this feature because one of the older operating systems, Windows* 98 (Win98), was not designed to deal with hard drives when the OS is first powering up. This situation, which is also seen on some portable PCs, can cause significant problems.

Normally Win98 doesn't deal with the hard drive until after POST. However, in S3, the hard disk drive loses power, just like other system components. During wake, drives receive power and go through their start-up sequence. During this wake, some hard drives return a 00 status, indicating they are busy. This is a valid status, but the OS interpretation of that status is wrong: the OS just doesn't recognize the code. Specifically, Win98 assumes that the 00 status indicates that the hard disk drive is not busy. Win98 then starts sending commands to the hard drive, a process that could prevent the system from fully resuming from S3.

For customers who want Win98 as their operating system, Compaq provides a hard drive reset to prevent the busy-status problem. By configuring BIOS to do a hard-disk reset upon S3 restore, the Deskpro EN ensures that Win98 doesn't get a chance to talk to the hard drives until the drives are returning a status that the OS understands.

USB Device Power

As human interface devices, USB (Universal Serial Bus) mice and keyboards should be able to wake the PC from the S3 STR sleep state. However, some OEMs do not provide power to the USB ports during S3 STR. In these systems, none of the wake-capable USB devices can wake the PC from the S3 sleep state.

Keyboards and mice must receive power in order to generate the wake signal. Compaq's approach is that all USB ports should receive trickle power in the sleep state, and so all USB ports on the Deskpro EN receive power while in S3 STR. This eliminates the need to have users press the power button to wake the PC, and PS/2 allows them to press a key on the keyboard to restore volatile system context.

PS/2 Mouse and Keyboard Wake Up

Some PCs do not provide power to the PS/2 ports in S3 Suspend-to-RAM. This means that users must press the power button to resume the system after the system enters the suspended (S3) sleep state. However, especially for users who left their Microsoft Excel, Microsoft Word, or other application open, pressing the power button can seem like a dangerous act.

Most users prefer to use a more seemingly benign keystroke or mouse movement to wake the PC from a sleep state. Because of this, Compaq put the PS/2 devices on the auxiliary well. This costs very little in terms of power usage; the Deskpro EN still draws 2 watts or less in the S3 state.

Design issues with other sleep states (S4 and S5) also offer opportunities for using S3 implementations to differentiate products. For example, electrically, the shutdown state (S5) is no different from hibernation (S4). Main power is cut off, and there is very little power on the resume well. In S5, the system is not concerned with wake-capable devices. And, in some systems, there may be almost no power used at all. This also means that PS/2 devices (such as keyboards and mice) are not powered in the S4, or hibernation, sleep state.

In terms of software, S5 is different from S4 because the OS will save memory to the hard drive before it shuts the PC down. When powering back up, the PC goes through POST again. Most BIOSes don't know the difference between power-down and hibernation. However, the OS does know, and the OS restores the appropriate information to the system.

To users, S5 (shutdown) means that they should be able to cold-plug devices (such as memory, PCI cards, and PS/2 devices) into the system. To users, PS/2 components should not receive power during the shutdown (S5) periods. However, because of the electrical characteristics of the system, PS/2 devices that receive power in S4 also automatically receive power during S5. In such a situation, the only way to cut off power to PS/2 devices in S5 is to unplug the system from the external power source (UPS, wall power, etc.) so that the devices can be cold-plugged. Otherwise, PS/2 devices plugged in during shutdown are actually hot-plugged, and can be damaged.

Consequently, PS/2 devices are not powered in S4. However, the Deskpro EN resolves this issue by providing power to all PS/2 devices during S3, so that users can easily press a key or move the mouse to restore the previous volatile system context.

Summary
Compaq's Deskpro EN series shows how one manufacturer is using the new S3 technology to differentiate its products from other platforms.

To provide a robust solution for any platform implementing Instantly Available PC technology, developers should consider solutions that involve BIOS. Compaq's approach involves custom designs for components such as power supplies. Compaq's approach also involves effective BIOS setup features that allow users to compensate for the current lack of power budgeting features in the OS, and for applications that are not yet completely compatible with S3.

More Info
"ENERGY STAR and Instantly Available PCs" by Jill Abelson, *Intel Developer Update* magazine, July 2000.

"Power Budgeting Protects the Instantly Available PC" by Patrick Bohart and Ram Chary, *Platform Solutions News*, April 1999.

Author Bio
Louis Hobson has been with Compaq for 10 years. He has worked on such projects as test tools, drivers, and BIOS development, and is currently with the Compaq Commercial Desktops Division. Previously, he worked in digital signal processing in the oil and gas exploration industry. Louis holds four patents, with others pending. He holds a B.A. in economics from Southwestern University, and an M.S. in mathematics from Memphis State University.

Initiatives and Technologies

Real-Time 1394b Data Transfer for Consumer Electronics

Steve Bard
Senior Staff Engineer, Platform Architecture
Mobile Architecture Lab
Intel Corporation

Overview

Attempts to network consumer electronics (TVs, VCRs, DVDs, etc.) with desktop PCs have resulted in complex and expensive cabling nightmares. The latest solution to the interconnect cabling problem is called IEEE 1394, more popularly known by names such as FireWire*, iLink*, and Lynx*.

In March 2000, the IEEE P1394b Working Group released and made available online the IEEE 1394b Revision 1.0 Amendment Draft Standard, an unapproved draft (meaning it is subject to change). Simultaneously, the 1394 Trade Association, at their Developer's Conference 2000, released a white paper and a presentation on 1394b. This paper discusses the most cost-effective and "future proof" methods of merging 1394 consumer electronics (CE) interconnectivity with desktop and mobile or notebook personal computers (PCs). Based on the new draft standard, the white paper recommends that PC OEMs planning to implement 1394 into their products do so by incorporating the new, extremely fast 1394b serial-bus PIL/FOP interface that is described in the IEEE P1394b Draft Specification.

There are two main considerations for adding 1394b (versus 1394-1995 or 1394a-2000) capability to desktop and notebook PCs: power consumption and cost. New technology must be extremely cost sensitive because of stiff competition in the consumer PC market. And, with the current Environmental Protection Agency (EPA) ENERGY STAR requirements, PC manufacturers must comply with increasingly strict guidelines for energy consumption.

The mobile PC (notebook) is even more constrained with regards to both cost and power. Because notebook PCs have limited-capacity batteries OEMs must carefully consider the power consumption levels of each subsystem, including external interconnects.

For both desktop and notebook PCs, long-term solutions to the current interconnect mess must resolve issues of cost, power consumption, and constantly advancing technology.

What Is 1394b?

While 1394 is the premier CE interconnect for devices requiring real-time, high-speed data transfer, 1394b is the best instantiation of 1394. This interface promises to deliver a single interconnect cable for all audio and video devices, including isochronous (video) mass-storage devices (iHDD, CD-ROM, etc.), TV, radio receiver, CD player, VCR, DVD, cable set-top box, camcorders, and digital still-image cameras. For some devices, such as Digital VCRs and camcorders, 1394b allows OEMs to exploit host-based (PC-centric) application processing for lower cost CE devices. This relieves OEMs from having to incorporate into the CE device the cost associated with application processing.

The 1394b interface supports high transfer speeds. For example, some instantiations of the interface support data transfer rates from 800 Mbits/second to 1.6 Gbits/second. These instantiations of 1394b can also be modified to support future needs for even higher data transfer rates.

The 1394b interface also allows for transmission distances up to 100m—longer than either the earlier 1394 or amended 1394a interfaces.

Finally, the 1394b interface allows OEMs to use the most appropriate transfer media for the transmission distance and data-transfer speed for the application. Traditionally, copper cable has been used in all implementations of 1394 up until 1394b. However, although developers can use a single copper cable for speeds up to 400 Mbits/second, they might not want to use copper cable in certain situations. The 1394b interface offers both the new concept of long-haul use, and the opportunity to use different types of media (for higher speeds over longer distances), depending on the performance requirements of the system.

Data-transfer media can include:

- Plastic optical fiber; transmission distance up to 50m, data transfer up to 400 Mbits/second.
- Multimode glass optical fiber; transmission distance up to 100m, data transfer up to 1.6 Gbits/second.
- Standard 1394 copper cable; transmission distance up to 4.5m, data transfer up to 400 Mbits/second. Longer cable lengths require that the construction characteristics of the cable change to match 1394 electrical/performance requirements.
- UTP (untwisted pair), also known as CAT5 cable; transmission distance up to 100m, data transfer up to 100 Mbits/second.

(The IEEE 1394b standard does not cover wireless technology; however, wireless transmission distances vary widely depending upon current, nonstandard implementations of the technology.)

The 1394b interface resolves the complicated wiring issues that have traditionally arisen when users tried to connect CE devices with PCs. Instead of offering a confusing variety of cables for traditional networked systems, 1394b provides ease of use and high bandwidth for moderate cost. With 1394b, consumers can connect their own PC platforms to take advantage of the powerful features and enhanced capabilities of modern CE devices.

The 1394b draft standard is an amendment to IEEE 1394-1995 Standard for a High-Performance Serial Bus. This new draft standard is at revision level 1.0, and has been submitted to an IEEE ballot body for sponsor ballot. An unapproved IEEE 1394b draft (revision 1.0), subject to change, is available for public review.

USB versus 1394b

The 1394 and USB technologies are complementary. USB 1.1 is a low-cost, medium bandwidth connection for keyboards, mice, low-to-medium resolution scanners and printers, phones, digital still cameras, and similar I/O devices. The 1394 interconnect is a moderate-cost, high-speed bus designed to bring high-speed/high-resolution printers and scanners, isochronous (video) mass-storage devices (iHDD), DVD, camcorders, and digital cameras into the PC-CE fold.

The two technologies (USB and 1394b) will exist together on future PC platforms. This is because low-speed USB devices such as keyboards and mice do not have the same requirements for high-speed interconnects as digital cameras, DVD, and other CE devices. Keyboards, mice, and other such devices will likely continue to use USB interconnects.

Cost difference between USB and 1394 is one reason the newer high-speed 1394 serial interface is expected to appear, primarily, in devices that require large and very high-speed throughput. As 1394b becomes more common, costs of the interconnect will drop.

PC-CE Topology

The big question now is where does the PC fit into the topology of this interface? PCs have continued to decline in cost and improve in performance and versatility. This makes the PC the favored platform and tool for networking consumer electronic (CE) devices. With the high-speed, serial 1394b interface, users will be able to easily connect not only traditional PC peripheral devices, but also CE devices to a PC. This means that the PC would be able to effectively network not only standard data storage devices (HDD, DVD-ROM, CD-RW, etc.), scanners, and printers but also CE devices that provide a 1394 interface (audio devices, camcorders, VCRs, digital video-editing devices, digital cameras, TVs, and so on).

Connecting PC and CE devices enhances the capabilities of both the PC and the CE devices. For example, networking PC and CE devices improves the user's ability to do computer animation, Web content creation, and home entertainment. Video editing, titling, music editing, audio video management and/or sharing, M-mail (movie mail), device control/management... these enhanced capabilities are a significant and compelling reason to add the PC system to a network of 1394 devices.

PCs can provide a central control and configuration point for A/V devices, allowing consumers to network and control their computer applications, video-editing applications, and home-entertainment devices from a single location. However, in a 1394 network, any device can be connected to any other device. In other words, 1394 is a peer-to-peer interconnect.

For example, a camcorder, a digital TV, and a digital VCR can be connected to a PC via 1394; or, the camcorder, digital TV, and the digital VCR can be connected directly together. The new 1394b draft standard even resolves the old problem of loops. Bus initialization problems due to the circular connection of 1394 devices is automatically resolved in a network of 1394b devices.

Implementation

IEEE 1394b PIL/FOP has several advantages over previous versions of the interconnect standard. The latest draft:

- Describes a backward-compatible architecture that allows users to connect and operate all 1394, 1394a, and 1394b devices.
- Suggests the lowest overall cost for instantiation of a 1394 system in a PC.
- Provides the best energy conservation mechanisms available for 1394 implementations.

By choosing the right instantiation of 1394b, manufacturers can achieve maximum, real-time data throughput for the lowest cost and still provide significant product differentiation. With 1394b, manufacturers can differentiate products based on cost, versatility, performance, energy efficiency, and ease of use. One instantiation in particular, mixed-mode FOP expansion ports, appears to be the lowest cost, most versatile, future-proof realization of this interconnect technology.

All implementations of 1394 require simple Intellectual Property (IP) licensing for 25 cents ($0.25 US) per system. The IP license provides for any number of 1394 devices in a system and any number of ports on each device.

Specific instantiations of IEEE 1394b architecture repartition the cost model of implementing 1394b on the PC based on various market requirements. For example, backward compatibility with older implementations of 1394 devices is important. However, legacy removal is also critical for reducing system cost. For this reason, PC OEMs should carefully consider not implementing previous generation 1394 hardware. Instead, PC OEMs should provide connection to previous generation of 1394 devices via a translator interconnect cable. Such a translator cable would plug into a standard 1394b socket on one end and on the other end, into either a 1394-1995 or a 1394a-2000 device. Users who need to connect their PCs to older, 1394- or 1394a-compatible devices can do so with this simple, after-market translator interconnect cable.

Another instantiation—the serial PHY/Link interface (PIL/FOP)—is more cost-effective and flexible than the traditional, parallel PHY/Link interface. The PHY/Link interface provides a bit of insulation against ever-advancing 1394 technology. Using this instantiation means that, as technology continues to move forward, OEMs may not have to replace both the link component and the PHY component when porting or designing new versions of their products. Only the PHY component would need to be updated when a new CE device is designed.

Serial versus Parallel

The 1394b serial link interface gives OEMs an opportunity to integrate a 1394 Open Host Controller Interface (OHCI) 1.1 register-compatible link with a single port beta-mode PHY. This eliminates the major issues of integrating the analog circuitry found in 1394a PHY device. The integration creates a PHY-integrated-Link device, or PIL. The PIL has one exposed, fully operational, beta-mode PHY port. This port may be connected to a standard beta socket or to a special PHY called a Fan-Out-PHY (FOP).

When connected to a FOP, the PIL/FOP connection may serve as a serial PHY/Link interface. The interconnection between a PIL and a FOP does not have to exhibit serial PHY/Link interface behavior. However, when connected as a serial PHY/Link interface, the PIL and the FOP behave as a single 1394 node. In other words, they collectively occupy a single node identifier (a NODE ID). This is the preferred approach.

When not connected as a serial interface, the PIL has a 1394 NODE-ID distinct from the FOP NODE-ID. (This article assumes the interconnection of the PIL and FOP to be that of a serialized PHY/Link interface.)

Parallel Interface

The IEEE 1394b draft standard allows 1394b to be implemented with the more tradition parallel PHY/Link interface. This interface consists of a minimum of 16 signals that interconnect the Link with the PHY:

D0	D4	CTL0	Pint
D1	D5	CTL1	LClk
D2	D6	Lreq	LPS
D3	D7	PClk	LinkOn

One of the drawbacks of this approach is that the wide bus between the Link and PHY increases the cost of both components. It also causes printed circuit board (PCB) routing challenges. Another limitation is the maximum transaction speed. For a parallel PHY/Link interface, the maximum transaction speed is only 800 Mbits/second.

The parallel PHY/Link interface must be engineered for an impedance-controlled transmission. And designers may not use differential signaling to assist the transmission. This is because the trace to ground impedance has to be carefully controlled—something especially difficult when going through vias.

Although engineering through vias is the most difficult challenge, there are other disadvantages to the parallel PHY/Link interface. The skew between the clock signal and any of the eight data signals must be small. Also, the lines are signal-reversing in nature, and so engineers must incorporate bidirectional termination in their designs.

Serial Interface

The 1394b draft standard introduces a new model (the PIL) for providing a serial interface between the link component and the physical layer (PHY). This serial PIL interface resolves many of the issues of the parallel PHY/Link interface.

For example, the serial interface reduces the number of pins (signals) required between the PHY and the Link from between 16 and 22 pins to just 5 pins. By combining the chips and reducing the number of pins required, the serial interface also reduces the PCB space required for implementing the standard. Also, the signal lines on the PIL port are unidirectional, meaning that there is no signal reversing, thus eliminating any need for bidirectional signal-line termination. Another advantage is that, when required, DC isolation can be achieved with simple capacitive isolation—a benefit of DC-balanced signaling.

System design can also be simplified with the 1394b Serial PHY/Link Interface. This is because the serial PHY/Link interface between a PIL and a FOP allows the two components to be located in different parts of the system. For example, the FOP can be near the connector, while the PIL is near the processor and memory.

The PIL/FOP model enables platform differentiation via the FOP while using a standard base chip (the PIL). The FOP is available in various combinations of data strobe (DS), bilingual (both DS and beta-mode capable), beta mode, and any mixture of those technologies.

The serial PHY/Link interface also helps designers in laying out PCBs. Impedance-controlled transmission lines are still required for the differential signal pair. And while the differential skew between the propagation delays on the two traces must remain small—an engineering challenge that is still tricky through vias—the signal direction of one data pair is unrelated to the signal direction of the other pair. Each pair of signals can take a different route on the PCB. In other words, they can have different trace lengths.

The 1394b PIL architecture is also completely digital: there is no common-mode signaling. With a lack of analog components in the architecture, developers have many choices for silicon development processes for PIL design. This is of particular significance when considering the ability to voltage-scale the technology. Choice in the most effective voltage-scaling method translates to significant power consumption conservation (I^2R loss).

The serial PHY/Link interface comprises five signals: two Differential Transmission signal pairs (a plus and minus line for each) and one LinkOn signal. You can download the signal definitions and descriptions from the IEEE P1394b Draft Standard, revision 1.0, which is freely available to developers.

FOP Configurations

The Fan-out-PHY (FOP) is available in a variety of configurations. The most popular configurations of the PIL ports are:

- PIL Port with 1394b data-strobe expansion ports
- PIL Port with only 1394b beta-mode expansion ports
- PIL Port with 1394 bilingual expansion ports
- PIL Port with a mixture of 1394 expansion ports

Ports

Data-Strobe Ports

This FOP configuration is a PIL port with 1394b data-strobe expansion ports (a border node). The most simple and logical choice for early instantiations of 1394 in a PC system would make use of a border-node FOP. A border node is any node (such as a FOP) that creates a signal interconnect between a 1394b beta-mode connection and a 1394a-2000 or 1394-1995 data-strobe mode connection.

For example, a FOP would interface to the PIL using a single beta-mode port. Additional ports would provide connections only for data-strobe mode. This is likely the least costly PHY to use with 1394 technology. This device should be in the median cost category.

This FOP configuration is limited to transaction data speeds between 100 Mbits/second and 400 Mbits/second.

Beta-Mode Ports

This FOP configuration provides the PIL port with only 1394b beta-mode expansion ports. This FOP is currently available and fully supported in the 1394b draft standard. It is also likely to be the PHY interconnect of the future, when all 1394 devices have moved to 1394b technology. However, there may be uses for this PHY that have not been realized by those who have developed the architecture to support it. As with other future-proof technologies, the market will dictate the full range of its uses.

Perhaps the most obvious use of this device would be as a long-haul 1394 signal repeater. In this configuration, the FOP would bridge one cluster of 1394a devices to another cluster of 1394a devices over a long distance, such as 100 meters. The FOP would connect the two clusters through glass optical-fiber media and would communicate in beta mode.

This configuration differs from the first FOP instantiation (PIL port with 1394b data-strobe expansion ports) only in its expansion-port signaling capability. The first instantiation uses data-strobe mode only; this FOP instantiation uses beta mode only.

This FOP configuration supports the entire range of transaction speeds currently available: 100 Mbits/second through 1,600 Mbits/second. And, the 1394b architecture has also laid the foundation for this FOP to handle speeds of 3,200 Mbits/second and, perhaps, beyond.

Bilingual Ports

This third FOP configuration consists of 1394 bilingual expansion ports. This is the most versatile, but probably the most costly instantiation of the standard. Like other configurations, this one provides the PIL-port serial PHY/Link interface. However, the expansion ports of this FOP are dynamically able to determine their operating mode at the time they are connected to other devices.

In this FOP, each expansion port is able to operate fully in data-strobe mode or in beta mode. Each port is also capable of operating independently of other ports, or rather, in a mode different from another port on the FOP.

For example, one port may be connected to a digital camcorder that has a 1394a interface. A second port may be connected to a device that provides a 1394b beta-mode interface. In this configuration, the connection to the camcorder would be able to execute data-strobe signaling, while the connection to the 1394b device would be able to execute beta-mode signaling.

Mixture of Ports

This FOP is a PIL port with a mixture of 1394 expansion ports. This is probably the most cost-effective FOP configuration that is still reasonably versatile. It is similar to the other FOPs, but each expansion port is capable of operating only in one mode. Some expansion ports operate only in data-strobe mode; others signal only in beta mode.

This configuration is especially useful in a system architecture that requires a high-speed (or long-haul) connection, as well as a standard-speed connection to early-architecture 1394 devices. This type of FOP might be found in a PC that provides a long-haul A/V network connection between one cluster of CE devices and a PC system connected to another cluster of A/V CE devices.

FOP Locations

A serialized PHY/Link interface between PIL and FOP creates many unique opportunities for product differentiation. This means that OEMs can differentiate, not only by cost and ease of use, but by location of FOP and by the number of interconnect options provided. For example, the FOP can be located in the PC system itself, or for notebook PCs, in the AC power adapter.

- *FOP in the PC system*. When the FOP is located in the PC system, it provides the standard CE device interconnect. However, when in the PC, the FOP can also provide a high-speed 1394b interface to an 1394b internal device or a long-haul A/V network connection.
- *FOP in the AC power adapter*. Notebook PCs provide OEMs a unique opportunity to locate the FOP inside the AC power adapters. Designers could even replace the current power plug on the notebook with a beta-only socket. The connector end of the adapter's tethered cord would then be a 1394b beta plug. An AC power adapter that contains a FOP can provide users with a number of 1394 interconnect options. See Figure 1.

FOP in the AC Power Adapter

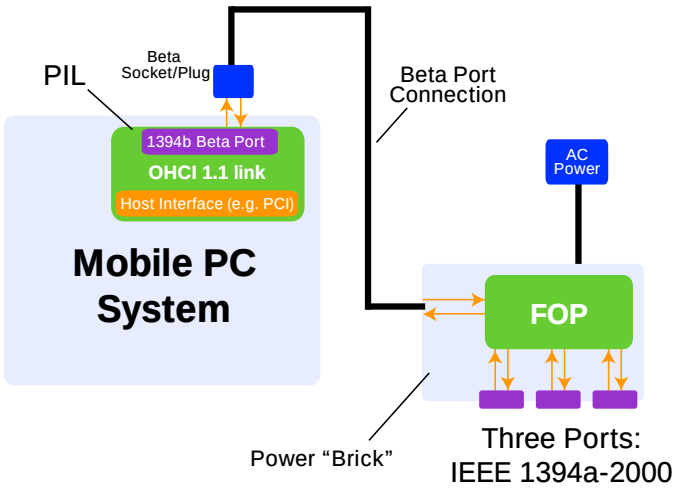

One question raised by the second approach is how can the notebook's beta-mode port be useful without the AC adapter? The answer lies in the beauty of the PIL concept.

Remember that the PIL port is a fully functional 1394 beta-mode port. Any 1394b device can connect directly to that port. The more important issue is regarding early 1394 devices that do not have a 1394b interface. For example, such devices could include camcorders, 1394a set-top boxes, and other devices that may provide broadband Internet connectivity. The question here is how do those earlier 1394 and 1394a devices connect?

With PIL, these devices can make use of a two-port, border-node FOP. One port is a beta mode, while the other port is a data-strobe mode. Such a FOP is extremely simple and very small. Theoretically, such a FOP could be embedded in the 1394 cable or in the plug shell. This means that 1394 and 1394a devices can connect to the single PIL beta-mode port via a simple cable containing a two-port border node FOP. See Figure 2.

Simple Cable with Two-Port Border-Node FOP

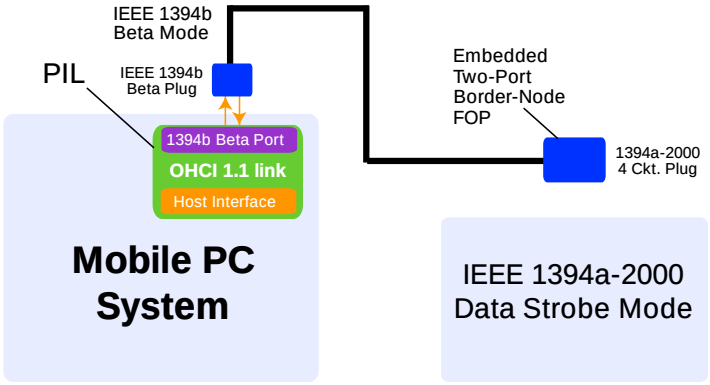

A two-port, border-node FOP cable would have uses other than for notebook PCs. To some extent, there are no restrictions for this type of cable use. Any interconnect requiring a translation from beta mode to data-strobe mode could make use of this cable architecture.

Summary

IEEE 1394b promises to deliver a single interconnect standard for all audio and video appliances. This includes TV, radio receiver, CD player, DVD, VHS, cable set-top box, etc. However, modern video editing, titling, and display features require more bandwidth and power from the PC than traditional audio and visual (A/V) devices. Today's camcorders, VCRs, digital cameras, and other A/V devices require a common infrastructure to take advantage of the host processing power of the home PC. With its flexibility and future-proof requirements, IEEE 1394b may become just that home network infrastructure for all A/V consumer needs.

Consumer PCs that incorporate 1394 capability should be based on the 1394b PIL architecture. This is the preferred 1394 building block for modern PCs. However, there are many PIL configurations that OEMs can use to achieve maximum bandwidth and the highest data-transfer rate. The many available configurations will allow OEMs to differentiate their products on a variety of levels, including cost, versatility, performance, and ease of use.

More Info

For a copy of the white paper released by the 1394 Trade Association at their Developer's Conference 2000, contact:
James Snider, Chair
1394 Trade Association
Regency Plaza, Suite 350
2350 Mission College Blvd
Santa Clara, CA 95054
USA

+ 408-748-9416

For more information about how USB and 1394 topologies can work together, refer to the *Intel Platform Solutions News* article, "USB and 1394 Living Together in Harmony" available through the *Intel Developer Update* magazine archives.Developers can study the unapproved (and subject to change) IEEE 1394b Revision 1.0 Amendment Draft Standard online. This draft contains the full specification for 1394b PIL/FOP and all additional 1394b technical details.

Contact the 1394 Licensing Authority (1394LA) for specific implementation costs.

Author Bio

Steve Bard is a senior staff engineer in the Platform Architecture Group of Intel's Mobile Architecture Lab. Active in the development of 1394 technologies for the past three years, Steve was instrumental in driving energy conservation mechanisms–specifically, suspend/resume and standby/restore–into the 1394 standard. Steve is an editor and major contributor to the 1394 TA Cable Power Distribution specification and Suspend/Resume Implementation Guidelines specification. He is also secretary to the IEEE P1394b Working Group.

PXE Manageability Technology for EFI

Mike Henry
Project Manager
Intel Architecture Lab
Intel Corporation

Overview

PXE (Preboot Execution Environment) is an existing open industry specification for enterprise network clients to automatically download software images and configuration parameters. A subset of Wired for Management (WfM) Baseline 2.0, PXE 2.0 is part of the soon to be released PC 2001 System Design Guide. Essentially all desktop PCs that feature an integrated network interface include support for PXE remote boot. Many ISVs, including Computer Associates, IBM, On Technology, Altiris, and Rembo Technology, have implemented PXE boot servers as part of their Microsoft Windows* server products. PXE boot service is a standard part of Windows* 2000 Remote Installation Service (RIS) and a standard part of Red Hat Linux* 6.2 server.

The EFI (Extended Firmware Interface) specification defines a new model for the interface between operating systems and platform firmware. The interface consists of data tables that contain platform-related information, plus boot and run-time service calls that are available to the operating system and its loader. Together, these provide a standard environment for booting an operating system and running preboot applications.

The EFI specification is primarily intended for the next generation of IA-32 and IA-64 Intel® Architecture-based computers, and is an outgrowth of the "Intel Boot Initiative" (IBI) program that began in 1998.

PXE 32/64 is now a standard component of EFI.

PXE Aids Manageability

The technology vision behind PXE technology is to enable the network interface as a boot device for Intel Architecture platforms, as common and familiar as a hard drive or CD-ROM. PXE technology enhances the manageability of networked client machines in several ways:

- *Remote new system setup*. If the client does not have an OS installed on its hard drive or has no hard drive at all, downloading a Network Bootstrap Program (NBP) from a server can automate operating system installation and other configuration steps.
- *Remote emergency boot*. If the client machine fails to boot due to a hardware or software failure, downloading an executable image from a server can provide the client with a specific executable that enables remote problem notification and diagnosis.
- *Remote network boot*. Where it is desirable to centrally administer the complete operating environment for a client, the client can download its system software image from the server in the course of normal operation.

PXE Primer

PXE uses the DHCP (Dynamic Host Configuration Protocol) and TFTP (Trivial File Transfer Protocol). DHCP is an Internet protocol defined by the Internet Engineering Task Force (IETF) to dynamically provide communications-related configuration values such as network addresses to network client computers at boot time. TFTP is an Internet protocol defined by the IETF to enable the transmission of files across the Internet.

In addition to the normal use of DHCP, PXE embodies two key technologies:

- *A boot server discovery protocol* for the client to locate an instance of a particular type of boot server, and request the downloading of a Network Bootstrap Program (NBP) from this boot server.
- *A set of APIs* available in the system's preboot firmware that constitute a consistent set of services that can be employed by the NBP or the BIOS.

With these capabilities, a new network client machine can enter a heterogeneous network, acquire a network address for itself from a DHCP server, and then download a Network Bootstrap Program (NBP) to set itself up or to use as its native operating environment. The PXE protocol defines methods commonly used in the industry that ensures this process can be completed across a wide variety of client platforms served by a variety of merchant PXE Boot Servers, and ensures that IT managers may centrally define and manage this network-based booting process for individual clients.

A New PXE for EFI

PXE was first introduced as part of the Wired for Management 1.0 specification in August 1997. It was upgraded to PXE 2.0 in the Wired for Management 2.0 specification in December 1998. PXE 2.0 was designed as an "option ROM" for the standard PC-16 bit BIOS. As a result, PXE 2.0 runs in real mode and provides 16-bit APIs.

EFI presents a protected mode 32- or 64-bit interface to the booting OS (unlike the standard PC 16-bit BIOS, which presents a 16-bit real mode interface). To include PXE in EFI required updating the APIs to 64 bits and rewriting the PXE core to a protected mode model. More than this, it required implementing PXE as an integral element of EFI.

PXE for EFI: Implementation

In keeping with the overall design philosophy for EFI, as much of PXE as possible was designed as an EFI protocol. As a result, the "PXE 32/64 BaseCode" is common to all network interface hardware. Network interface vendors are no longer responsible for implementing this portion of PXE functionality as it is now built in to EFI.

Under EFI, the PXE APIs to the network interface hardware (collectively known as the Universal Network Driver Interface, or UNDI) are now 64-bit APIs and have been implemented for both IA-32 and IA-64 instruction set architectures. UNDI is written as an EFI device driver and is implemented as a subprotocol under EFI SNP (Simple Network Protocol). UNDI presents a flat model native addressing interface to the booting OS. The new UNDI is faster and, in the case of UNDI 32, smaller than the equivalent UNDI 16 module in PXE 2.0.

Finally, as Figure 1 shows, PXE 32/64 BaseCode supports existing PXE 2.0 option ROMs. This approach provides for backward compatibility to preserve the functionality of existing PXE 2.0 enabled NICs.

PXE 32/64 BaseCode Support for PXE 2.0 Option ROMS

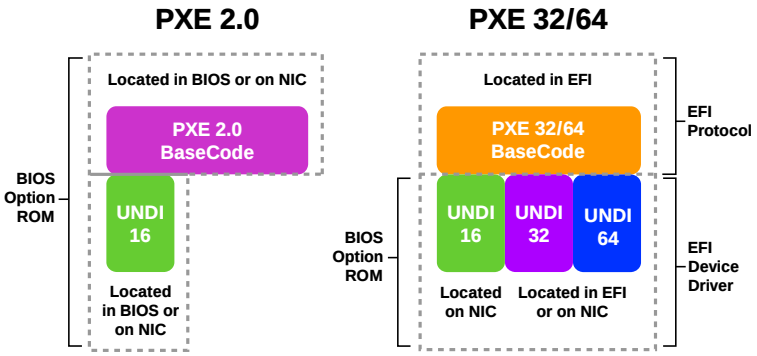

Summary

The EFI (Extended Firmware Interface) specification defines a new model for the interface between operating systems and platform firmware. The interface provides a standard environment for booting an operating system and running preboot applications. PXE 32/64 defines the standard network boot capability in EFI.

PXE (Preboot Execution Environment) makes it possible to remotely boot an OS from network clients regardless of the current content of the client's hard drive. By providing the open industry PXE specification, enabling tools, and interoperability testing, Intel is extending EFI to include the network as a standard boot device.

More Info

The specification for PXE 32/64 and the UNDI 32 and UNDI 64 API definitions are included in the EFI 1.0 specification.

You can download The Intel® Preboot Execution Environment (PXE) Software Development Kit (SDK) for Linux from the Wired for Management Tools section of the Intel Architecture Labs Web site. The SDK provides source code Linux developers can use to create PXE servers that are compatible with the EFI and the PXE specification. The SDK includes sample PXE and MTFTP daemon source code and build instructions, sample code for EFI Linux loaders for IA-32 and IA-64, and sample UNDI 32 and UNDI 64 implementations.

You can also download a copy of the PXE specification 2.0 in PDF format from the Intel Web site. This specification contains the wire protocol common to both PXE 2.0 and PXE 32/64.

You can download the EFI sample implementation and toolkit from the EFI Web site.

Author Bio

Mike Henry joined Intel in 1991. He is currently a project manager in Intel Architecture Laboratories (IAL) and is responsible for developing and executing PXE technology for Intel's Wired for Management initiative. Mike holds two patents and was a recipient of the Intel Achievement Award in 1998. He holds a B.S.E.E. and a B.S. in Psychology from the University of Washington.

Software

Remote Access to Pre-Release IA-64 Systems

Umesh Shah
Technical Marketing Manager
Intel Developer Services
Intel Corporation

Overview

Intel® Architecture 64 (IA-64) is a new, 64-bit architecture that enables highly efficient 64-bit computing. Because this technology is new, most software developers need some experience or training on an IA-64 system in order to take full advantage of this processor family's potential.

Public availability of IA-64 is not far away. With time running short, developers must have access to IA-64 systems now to make sure their applications are ready for market when IA-64 is released. However, many developers—especially those who work for small businesses—do not have access to IA-64 systems on which to develop and test their software. Developers who do have access to IA-64 systems also have to deal with maintenance and resource issues.

Intel is now offering developers free remote access to new technologies as early as possible. A new feature of Intel® Developer Services, Intel® Early Access Services (IEAS) gives developers Internet access to secure, maintenance-free, fully functional, multi-OS (operating system), software development environments for IA-64 applications. The service frees developers to concentrate on software applications instead of obtaining, maintaining, and upgrading systems. In addition, Intel Early Access Services provides new technologies both before and after they are publicly released.

With Intel Early Access Services, developers can check out new systems before purchasing them, perform both low- and high-level software test and development on cutting-edge systems, and ensure that their products are ready for market when new technologies are released.

Remote Development Environment

Intel engineers working across various geographic locations have been using remote computers for years to develop and test chips and software. Today's Intel Early Access Services is an extension of Intel's popular in-house, remote-access concept for the entire next-generation IA-64 processor family.

To help ensure end-to-end quality of service over the Internet Intel has eliminated network-related bottlenecks at the data center that houses the IA-64 systems for the service. The company estimated the number of users who would simultaneously access the data center at any one time, allocated appropriate bandwidth, then brought in a DS3 data-transfer pipe for handling it.
Intel provides for multiple levels of security and confidentiality through several methods, including user-application screening, authentication schemes, encryption, and static and live monitoring for unauthorized access and/or activity.

Other companies (HP, IBM, Microsoft, and others) are setting up their own IA-64 development labs. Intel's IA-64 data center offers several advantages over these labs, including multiple operating systems, flexible usage models, and additional security. Essentially, Intel's remote access IA-64 data center has all the attributes of an in-house data center.

Required Equipment

Any type of Internet hardware and applications can be used to access the remote IA-64 systems in IEAS data center. Customers can even use a dial-up modem, although such modems are rarely adequate for development purposes.

Intel recommends that developers use an Internet connection that provides at least 256 Kbits/second—anything that gets you on the Internet fast. The actual rate needed depends on the types of software development to be done. The two most preferred connection technologies are DSL (digital subscriber line) and ISDN (Integrated Services Digital Network).

Multiple Operating Systems

Intel currently offers IA-64 platforms with two operating systems:

- A 64-bit version of Windows* 2000
- HP UNIX* on IA-64

Intel is planning to add two more operating systems within the next few months:

- Linux*—by the beginning of October 2000
- AIX 5L (the IBM/Monterey OS)—by the end of December 2000

Shared Access

Intel offers customers two types of remote access to new technologies: shared access and exclusive access. The shared access model grants user privileges (not sys-admin privileges) to each user. This model restricts the system resources of each user, since each developer shares the system with other developers. However, this model does not restrict the amount of time a user can spend on the system. Resource limitations include hard disk space and system functionality.

Some examples of shared access use:

- Running sample code
- Training developers on the new architecture
- Understanding new tools (such as debuggers) that are available on the new platform
- Doing quick tests, for which you may not want to set up a complete system yourself

For example, a new version of an OS might have been released, and you now want to test your application on this new OS to see if the old bugs have been fixed. In this case, you could use shared access for a day to push your binaries on the shared system, execute the binaries, and then wipe them from the system.

Another benefit of this model is the way it ties in to Web tutorials available for new technology. For example, Intel's online tutorials for IA-64 are information-oriented, and include sample code. With shared access to a remote IA-64 system, developers can now run that sample code on a real system. This approach takes Web-based training to a new level.

Shared access over the Internet is inherently open. Protection from unauthorized access or actions is provided by the user's OS, the IA-64/OS platform, and via VLAN (virtual local area network) and other network security provided by the IEAS data center. Because of this moderate security level, Intel recommends the shared access model for purposes that are not sensitive to tight security.

Exclusive Access

The exclusive access usage model assigns an entire isolated network segment to a single developer or company. In this model, the developer (or developers from the same company) has full control over all system features and can use any and all 64-bit compute capabilities. The developer does not share the systems with other users, and all system resources (hard disk and functionality) are available to the assigned developer or company.

The exclusive access model offers three important benefits. One of the most obvious and important benefits is the additional level of security realized from isolation. Intel further protects an associated Internet connection with encryption technology.

Another important benefit of the exclusive access model is that, because of robust VLAN security features, Intel can also grant that developer full sys-admin privileges within the assigned network segment. This means developers can do high-end testing and development of applications for the new technologies. For example, sys-admin privileges are usually needed for testing and developing:

- File system drivers
- System software
- Kernel drivers
- Client-server applications
- Applications that require a lot of logic resources (memory and CPU power)

The third major benefit is that, within an isolated network segment, Intel can offer developers more than one platform on which to work. For example, by default, each isolated network segment includes both an IA-32 system and an IA-64 system. Theoretically, Intel can install up to six IA-64 systems or IA-32 systems on a single exclusive access network segment.

Exclusive access is granted for two weeks at a time. Developers who feel they need more time on the system can apply for additional time.

Multiple IA-32 and IA-64 Systems

Exclusive access users are assigned exclusive access to a private network segment that includes a minimum of one IA-32 and one IA-64 platform. This lets developers use cross-platform tools and set up a client-server development environment. Having remote access to multiple platforms gives customers a fully controllable, richer environment for software development.

For example, developers can configure their assigned IA-64 system as a server and the IA-32 as a client. They can then do client-server application development. If porting existing software to the IA-64 platform, they may have test scenarios where the test applications are driven from another system. In this case, the developers can configure the IA-32 system as the system that drives the tests for the 64-bit programs running on the IA-64 system.

Today's tools—the compiler and linker used to generate the binaries for the IA-64 system—are cross-platform tools. Developers run the compiler and the linker on an IA-32 system, and the 32-bit system generates an IA-64 binary. That IA-64 binary cannot be executed on the existing 32-bit systems. However, with remote access, early release program, 64-bit systems are available now for software development. This gives exclusive access users two options:

- Keep the source code on their local system at their own site, using the cross-platform tools on their IA-32 systems at their own site to create binaries. Then push the binaries to the IA-64 system and test the code on a remote access, 64-bit system (shared or exclusive access network segments). The disadvantage is that this forces developers to transfer large binary objects across the Internet connection, a process that can be very slow.
- Take the source code and move it to the remote IA-32 system on an isolated network segment (exclusive access). Then run the cross-platform tools remotely and transfer the resulting binaries to the remote IA-64 system on that same exclusive access network segment. This gives developers an extremely fast way to debug code for new technology.

Having multiple platforms available to customers is an important benefit of the exclusive access model for IA-64 software development. Along with multiple OSes and continuing availability of other new technologies, the early access program offers a unique and powerful opportunity for software developers across the industry.

Security and Confidentiality

Intel is as concerned about providing a secure and confidential environment as developers are in using one. For this reason, Intel provides both static and live security and auditing at the physical, network, and system levels. Security and confidentiality for the Intel Early Access Services' data center will be described in an upcoming article in this magazine. An overview is available now on the Intel Early Access Services Web site.

Who Qualifies?

Use of Intel Early Access Services is free of charge; however, potential users must apply and be approved for this remote access to new technologies. Citizens of controlled countries who are currently living and working in the U.S. may be limited to certain features of the service due to U.S. export control regulations.

Intel also looks closely at applications from Intel competitors or from companies in active collaboration with Intel competitors. Such applications may or may not be approved, depending on the individual situation and the purpose of the requested access.

Summary

Fully functional, maintenance-free IA-64 platforms are available now through Intel Early Access Services. Developers can start porting and developing their applications immediately with minimal effort—a huge advantage over waiting until IA-64 systems are publicly released.

Intel Early Access Services:

- Offers access to IA-64 systems before that technology is publicly released
- Offers IA-64 systems with a variety of operating systems
- Eliminates the need to buy, maintain, and upgrade a system yourself
- Allocates adequate bandwidth for most developer needs
- Provides high-speed, real-time data transfer using state-of-the-art Internet technology
- Offers technical assistance in setting up effective remote connections
- Offers different access models to meet the needs of different developers
- Provides security at a physical, network, and system level
- Provides effective and legally binding confidentiality

More Info

Developers can learn more about Intel Early Access Services at the IEAS Web site. To apply for the service, developers must first register at the Intel Developer Services home page. Intel will continue to offer developers the latest technologies through this remote access service.

For information about Intel Developer Services, see “Intel Developer Services Enables Better Solutions,” published in the July 2000 issue of Intel Developer Update Magazine.

Author Bio

Umesh Shah has been with Intel for nine years. He joined Intel as a VLSI chip designer for the i360SL and was member of a design team for i486SL processor, then worked in software design for the Mobile Companion (PDA). For seven of those years he has focused on Pentium® processor projects, including real-time address tracing, the Instruction Tracer (Pentium and P6 microarchitecture), the event-based profiler (similar to the VTune™ analyzer), and IA-64 server software enabling.

Currently, Umesh is a technical marketing manager with IAG/SEG/Developer Services and project manager for Intel Early Access Services. He received his B.S. in electrical engineering from BITS, Pilani, India; his M.S. in computer engineering is from Arizona State University.

—End of Intel Developer Update Magazine Issue 13—